AMENDMENTS TO THE SPECIFICATION:

Page 1, immediately preceding the paragraph commencing at line 3 ("The present invention relates to networks..."), insert the following heading and sub-heading:

BACKGROUND

1. Technical Field

Page 1, line 17: delete "Background" and insert the following sub-heading:

2. Related Art

Pages 2-3, bridging paragraph:

The following field in an IPv4 datagram header is a 16-bit "Checksum" field. Some values in [[a]] an IPv4 datagram header may change at each packet switch hop, so the checksum may need to be adjusted on its way through a network. The checksum is followed by 32-bit "Source Address" and [[a]] 32-bit "Destination Address" fields, respectively.

Page 5, line 12: delete "Summary of the Invention" and insert the following heading:

BRIEF SUMMARY

Pages 6-7, bridging paragraph:

According to the initially filed claims of this co-pending application, there is provided a data network comprising a provider node, a receiver node, and a plurality of intermediate nodes, the provider node being arranged to provide data to at least one of said intermediate nodes or to the receiver node, said intermediate nodes being arranged to receive data and forward data to at least one other intermediate node or to the receiver node, and the receiver node being arranged to receive data from at least one intermediate node or from the provider node; wherein:

said data comprises at least a part which relates to a path ~~characterisation~~ characterization metric;

said provider node is arranged to assign an initial condition to the path ~~characterisation~~ characterization metric in respect of data provided by it;

said intermediate nodes are arranged to update the condition of the path ~~characterisation~~ characterization metric in respect of data they forward;

said receiver node is arranged to make available for the provider node information indicative of a discrepancy between the condition of the path ~~characterisation~~ characterization metric in respect of data received by it and a predetermined target condition for the path ~~characterisation~~ characterization metric; and wherein

said provider node is arranged to assign a different initial condition to the path ~~characterisation~~ characterization metric in respect of subsequent data provided by it in the event that it receives information indicative of such a discrepancy from said receiver node.

Page 7, 1st–2nd full paragraphs:

Closely related to the above, the initially filed claims of this co-pending application also relate to a feedback node for enabling an initial condition to be assigned to a path ~~characterisation~~ characterization metric in respect of data to be forwarded through a data network, said data network comprising a provider node, a receiver node and a plurality of intermediate nodes, said data comprising at least a part which relates to a path ~~characterisation~~ characterization metric; said provider node being arranged to assign an initial condition to the path ~~characterisation~~ characterization metric in respect of data, and to provide said data to at least one of said intermediate nodes or to the receiver node; said intermediate nodes being arranged to receive data from said provider node or from one or more other intermediate nodes, to update a condition of the path ~~characterisation~~ characterization metric in respect of data received by them, and to forward data to at least one other intermediate node or to the receiver node; and said receiver node being arranged to receive data from at least one intermediate node or from the provider node, and to make available for the feedback node information

relating to the path ~~characterisation~~ characterization metric in respect of data received by it; wherein

the feedback node is arranged to enable a different initial condition to be assigned to the path ~~characterisation~~ characterization metric in respect of subsequent data provided by the provider node in the event that said feedback node receives information indicative of a discrepancy between a predetermined target condition for the path ~~characterisation~~ characterization metric and the condition of the path ~~characterisation~~ characterization metric in respect of previous data received by said receiver node.

In general, it will be understood that the variable "condition" and the "predetermined target condition" for a path ~~characterisation~~ characterization metric will usually be values, examples of which are provided in detail below. It is foreseeable, however, that certain embodiments of the invention may instead use types of conditions that are not themselves values, such as, for example, the amplitude or phase of an optical signal in an optical network.

Pages 7-8, bridging paragraph:

Embodiments of the above may allow the following to be achieved:

1) Provision of path ~~characterisation~~ characterization information to nodes in a network, said information relating to any of a variety of possible characteristics of the

path or paths downstream of the node in question. To achieve this, there may be no need for upstream traffic beyond that being fed back end-to-end from a destination of data to the appropriate source. This is particularly useful where routes are asymmetric, particularly where it is not possible to send data upstream over certain unidirectional links (e.g. satellite links). But it is also useful if the available capacity can be increased by removing the overhead of routing information.

2) Ensuring that information such as the above may be proofed against falsification for the gain of an individual controlling any intermediate or end node.

Page 8, paragraph commencing at line 13:

It will be evident that while the path ~~characterisation~~ characterization metric may in effect "travel" through the network with the item of data to which it relates, for example, in the header of a data packet, according to a new version of an Internet Protocol, for example, this is not necessarily the case. A path through a data network may be no more than a virtual data channel, and there is no need to restrict the "location" of any path ~~characterisation~~ characterization information (to the extent that information has a location at all) to being within that channel. For instance, many network technologies separate control information from the data to which it refers. Control information, such as that ~~characterising~~ characterizing paths, is carried in separate messages using separate protocols, that refer to the relevant data channel

that they ~~characterise~~ characterize. Sometimes control information is even carried on separate physical links between control equipment that is distinct from data forwarding equipment. More commonly, control information is carried in separate virtual circuits over the same physical links. For this reason, variants of the invention as set out above may perform the path ~~characterisation~~ characterization steps remote from the network, rather than within the network.

Pages 8-9, bridging paragraph:

It will be noted that information corresponding to that which can be made available to intermediate nodes according to the above method, which allows for information relating to the downstream path to be deduced, can also be made available without assigning a different initial condition to further path ~~characterisation~~ characterization metrics provided that information relating to the difference between the eventual condition of a previous path ~~characterisation~~ characterization metric and said predetermined target condition is made available to the intermediate node in addition to information indicative of the updated condition of further path ~~characterisation~~ characterization metrics. Using these two pieces of information, information relating to the downstream path can similarly be deduced in respect of a particular intermediate node.

Page 10, 1st-2nd paragraphs:

The invention of the present application relates to the treatment of data, such as the routing of data in a data network. Path ~~characterisation~~ characterization information such as that derived according to methods referred to above is capable of being used by intermediate nodes in a network when making routing or other decisions. Such decisions may be based on more directly relevant, useful and up-to-date information than has previously been possible, provided that such intermediate nodes are capable of deriving appropriate information from the path ~~characterisation~~ characterization information they receive.

Thus, according to the present exemplary embodiment ~~invention~~, there is provided an intermediate node for controlling the treatment of data in a data network, the data network comprising said intermediate node, at least one upstream node, and a plurality of downstream nodes, the or one of the upstream nodes being arranged to provide data to said intermediate node, the or one of the upstream nodes being arranged to provide path ~~characterisation~~ characterization information to said intermediate node, and said downstream nodes being arranged to receive data via paths downstream from the intermediate node; said intermediate node comprising:

means for receiving data from an upstream node;

means for receiving path ~~characterisation~~ characterization information from an upstream node, and for deriving therefrom information indicative of a characteristic of a path downstream of said intermediate node;

means arranged to select, in dependence on said information indicative of said characteristic of a downstream path, a preferred manner of treatment for data to be forwarded on a downstream path; and

means for forwarding data to a downstream node according to said preferred manner.

Pages 10-11, bridging paragraph:

Corresponding to this, there is also provided a method for controlling the treatment of data to be forwarded from an intermediate node in a data network, the data network comprising said intermediate node, at least one upstream node, and a plurality of downstream nodes, the or one of the upstream nodes being arranged to provide data to said intermediate node, the or one of the upstream nodes being arranged to provide path ~~characterisation~~ characterization information to said intermediate node, and said downstream nodes being arranged to receive data via paths downstream from the intermediate node; said method comprising the steps of:

receiving data from an upstream node;

receiving path ~~characterisation~~ characterization information from an upstream node, and deriving therefrom information indicative of a characteristic of a path downstream of said intermediate node;

selecting, on the basis of said information indicative of said characteristic of a downstream path, a preferred manner of treatment for data to be forwarded on a downstream path; and

forwarding data to a downstream node according to said preferred manner.

Pages 11-12, bridging paragraph:

In general, embodiments of the present invention are described with reference to data networks[[,]]; however, it will be noted that some embodiments of the invention may be applicable to other forms of network, such as workflow routing, electricity generation or even transport networks such as railway networks. However, the ~~principal~~ principle advantages of the invention are more evident in situations where it is problematic to provide immediate feedback for each message or event against the normal flow in the network along each link in order to keep each node in the network informed of the state of affairs downstream. Where the network carries non-information items (work, electrical current, cars, etc.), often it is not natural to be able to carry feedback backward along every link of the network, because feedback is often pure

information, which the network is not designed to carry. However, it may be sufficiently cost effective to arrange for items flowing forwards through the network to carry information even if they are not pure information themselves (e.g. cars), and for communications links to be strategically placed across inputs and outputs of the network to allow feedback to be returned to relevant inputs and re-inserted into the network. In cases such as these, embodiments of the invention might prove useful on a hop-by-hop basis back to the source. Given work-flows arrive much more slowly than packets, it may well be more efficient to send information directly back to the source of a workflow after each step of the process, than to piggy-back the information only on work-flows flowing forwards through the system. The distinguishing feature is that the "atoms" of messaging dealt with by a work-flow routing system are much larger than feedback to the source needs to be. This is a reason why feedback provided according to embodiments of the invention is particularly useful in a data network such as a packet network. It is advantageous to try to avoid sending feedback as often in the upstream direction as data is being sent in the downstream direction.

Page 12, 1st full paragraph:

Nonetheless, it will be apparent that embodiments of the invention could also be applied to connection-oriented networks with connections consisting of cells, frames or packets (e.g. ATM, Frame Relay, SDH, etc.). It could be applied to control of future all-

optical packet networks, which are hard to design because of the inability to buffer light packets arriving simultaneously at a switch so that they may be served sequentially. Although light cannot be stored without converting it to electrical information, it can be slowed down. Feedback provided according to embodiments of the invention could provide a mechanism to slow it down before it arrived at a contended switch output, as the inability to store light means the slow-down must be instigated in advance of it being needed, not once the light arrives at the output. Feedback provided according to embodiments of the invention is wholly applicable to routed overlay networks on the current Internet, such as those created in a peer-to-peer fashion like CAN, PASTRY, Chord and SWAN, described in the following publications:

Page 13, paragraph commencing at line 15:

Figure 4 is a graph illustrating the use of a path ~~characterisation~~ characterization metric based on an Explicit Congestion Level (ECL);

Page 13, line 19: delete "Description of Preferred Embodiments of the Invention" and insert the following heading:

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

Page 14, 1st full paragraph:

The provider node 22 and the receiver node 26 need not be the original source of the data or the eventual destination of the data. In this case, the originating source of the data is shown to be at node 20 which is outside the network 21, and the intended eventual destination of the data is shown as being at node 27, also outside the network 21. The only distinguishing features of a provider node and a receiver node relate to the fact that a receiver node sends feedback to a provider node which includes path ~~characterisation~~ characterization information.

Page 14, 3rd full paragraph:

In the IPv4 header, two fields are used to ~~characterise~~ characterize the path, the TTL and the ECN fields (certain options such as a "timestamp" option were also designed for this purpose). An embodiment of the invention that aimed to ~~characterise~~ characterize the path against hop count and congestion metrics may require modifications to the standards for handling IP headers. Therefore, the IP version field might be set to some future version, say eight. We will describe the embodiment using a new "Explicit Congestion Level"[["]] (ECL) field consisting of an 8 bit real number replacing the two bit ECN field (how this fits into the header need not concern us here). The TTL field could remain the same size, but both TTL and ECN fields will be used differently from their ~~standardised~~ standardized semantics in IPv4. As will be

understood from the explanation below, such an ECL field will be capable of providing path ~~characterisation~~ characterization information to any node, such path ~~characterisation~~ characterization information providing information from upstream of the node in question which is indicative of the amount of congestion likely to be experienced on a path downstream of the node in question by a data packet at the node in question.

Pages 14-15, bridging paragraph:

When providing a first data packet, the provider node 22 assigns values to various fields in a header associated with that data packet, which may include any or all of the fields explained above with reference to the Internet Protocol IPv4 with alterations similar to those just described. The provider node 22 assigns an initial value to what will be referred to as the "path ~~characterisation~~ characterization metric". As will be explained, the semantics of the path ~~characterisation~~ characterization metrics differ from those of the IPv4 header in a fundamental way, which is that the common reference level of the path ~~characterisation~~ characterization metric is arranged to sit at the receiver node 26, rather than the provider node 22.

Page 15, 1st full paragraph:

In order to explain this difference, reference will again be made briefly to the "time-to-live" (TTL) field in the Internet Protocol header. As explained earlier, this is

currently ~~initialised~~ initialized at the sender with a value of 255, and is decremented by every node that each packet traverses. Thus, at any node in the network, the difference (255-TTL) ~~characterises~~ characterizes the number of upstream hops that a packet has traversed. If the packet reaches its intended destination after 45 hops, the TTL value will have been decremented to 210, and will have served the purpose of indicating to intermediate nodes on the path that the packet had traversed no more than 45 hops. If, however, the packet was incorrectly routed and/or entered a loop such that it performed sufficient hops (i.e. 255 hops) for the TTL value to reach zero, this would indicate to a subsequent intermediate node that the packet could be discarded. In this event, an indication may be sent to the provider node that the packet failed to reach its destination, but subsequent packets would still be assigned an initial TTL value of 255.

Pages 15-16, bridging paragraph:

Contrary to this, a path ~~characterisation~~ characterization metric corresponding to the TTL field would be assigned an initial value by the provider node 22 such that if the packet traverses the same or a similar path on a subsequent occasion, and every intermediate node 24 on the path decrements it by one, it should end up at a predetermined common reference level of, for example, zero at the receiver. In order to achieve this, the receiver node 26 should feed back the difference between the actual received value of the path ~~characterisation~~ characterization metric, and the

predetermined common reference level at the receiver (e.g. zero) to the provider node 22. The provider node 22 can then adjust or correct the initial value of the path ~~characterisation~~ characterization metric in relation to future packets to the same destination so that packets should generally arrive at the receiver node 26 with a value of or near zero. It will be noted that the first packet sent, or other packets sent before any feedback is received are unlikely to hit the zero target, and may accordingly be flagged as "guess" packets. Once feedback relating to a "guess" packet has been received by the provider node and used to adjust or correct the initial value of the path ~~characterisation~~ characterization metric in relation to a subsequent packet, the value of the metric, as updated by subsequent intermediate nodes 24 in respect of hops traversed by the packet, will convey information to each subsequent intermediate node that relates to remaining number of hops to the destination, i.e. the "downstream path" in respect of node.

Page 16, 1st–2nd full paragraphs:

With this new arrangement, any node in the network (whether provider 22, intermediate 24 or receiver 26) can read the value of the path ~~characterisation~~ characterization metric in any "non-guess" packet as the predicted remaining number of hops to the destination, albeit one round trip time ago.

An important, if not fundamental advantage of using path ~~characterisation~~ characterization metrics such as the above in the above manner will now be explained with reference to the "routing" of data packets through a network. As will become apparent, embodiments of the present invention allow intermediate nodes, taking the role of Internet routers, for example, to make informed decisions with regard to the onward routing of packets they receive, based on information relating to the dynamic state of the downstream path to the destination (i.e. the path between the intermediate node in question and the intended receiver). They are able to do this without the need for upstream routing messages along the path in use other than those from the eventual receiver node back to the provider node. Previously, according to IPv4, routing messages have been passed upstream between intermediate nodes typically every 30 seconds. With use of a path characterisation metric as set out above, and without the need for such upstream routing messages, the changing state of the downstream path may be known almost continuously (albeit delayed by one round trip). Even at nodes where no data is currently destined for a particular destination, explicit additional routing messages need only be sent from nearby nodes that are being continuously updated. Thus, routing can continuously adapt and converge to downstream changes, without the need to wait for regular routing updates from the path in use. These advantages are applicable to improving routing convergence and efficiency in a variety of types of network, but this advantage is of particular relevance in relation to more dynamic

scenarios such as where there is network mobility or in an ad hoc network, or where a more dynamic metric such as congestion as well as more stable metrics such as hop count are used to ~~optimise~~ optimize routing.

Page 17, 1st – 2nd full paragraphs:

Referring to Figure 3, a simplified representation of a network is shown in order to illustrate how embodiments of the invention allow for the provision of path ~~characterisation~~ characterization information to nodes which allows them to make informed decisions relating to the routing of data through the network.

Figure 3 indicates how routing decisions may be made using path ~~characterisation~~ characterization information derived according to embodiments of the present invention. It shows how such path ~~characterisation~~ characterization information may be exploited to route information towards a receiver by the "best" possible route. The word "best" seems to imply that the choice is subjective, but the sense in which the route is seen as the best can be chosen by selecting a metric corresponding to any of a variety of categories. Depending on the category of metric used, "best" may thus correspond to "cheapest", "least-congested", "most direct", or "least propagation delay" etc., or even a weighted combination of these. In order to simplify the explanation, we will consider the routing of data based on just one type of path ~~characterisation~~ characterization metric, "propagation delay", for example. In this case, it will be

assumed that a "propagation delay" field exists in the data headers of the network protocol in use.

Pages 17-18, bridging paragraph:

In Figure 3, senders S1 to S4 (squares) represent possible provider nodes, which may be single hosts or other networks from which data may be sourced. ~~sourced.~~ Receiver R1 represents a receiver node. Routers RT1 to RT6 (large circles) represent possible intermediate nodes on the path from a sender to a receiver. Each interface of each router is shown (smaller circles) holding the link cost of its locally-connected downstream link Δm_i. Where two possible interfaces exist and a choice may need to be made between them, the link costs of the respective downstream links are shown in respective small circles. Using the example of propagation delay, this can be measured by a simple echo request along each link (whether wired or not) at boot, for example. For fixed links, a re-measuring of this delay may be triggered if the underlying logical link changes its topology, for example. For wireless links, it may be appropriate to measure propagation delay more regularly, depending on the likelihood of mobility.

Page 18, 1st–2nd full paragraphs:

As each router accepts data, it decrements the "propagation delay" field by the propagation delay Δm_i of the link the data was sent over. For simplicity, the target value

m_z for the "propagation delay" field will be taken to be zero in this example. Then according to embodiments of the present invention, after the first round-trip, further data packets flowing towards receiver R1 at every router may carry a "propagation delay to destination" (PDTD) value m_i in their headers which represents what the remaining delay to R1 was on the last round-trip. This is represented by the "in-data headers" (numbers in heads of large numbered arrows), and may be treated by routers as a Path ~~Characterisation~~ Characterization Metric (PCM). Routers may thus maintain the PDTD values for one, two, or more interfaces in their internal routing tables (see numbers inside large circles). Where two (or more) values are held, each router need only "advertise" its single "least-cost" or "best" route to its neighbors ~~neighbours~~, but where a router may itself need to make a choice between the different interfaces, it may do this at any time simply by comparing the "least-cost" route with the "next least cost" route, or (in other terms) by comparing the "best" route with the "next best" route, at any time.

Routing messages, using a protocol similar to the current "Routing Information Protocol" (RIP) and containing PDTD values for R1 may be sent regularly from routers outwards from R1, every 30 seconds for example, unless a change triggers an immediate message. These are shown as numbers inside black arrows. These routing messages may, however, be suppressed where data is flowing along a link towards R1, since path ~~characterisation~~ characterization information may then be provided instead according to the invention.

Page 19, 5th full paragraph:

The use of path ~~characterisation~~ characterization information together with some routing messages from the routers nearby allows the cost from the router to the destination to be discovered in almost real-time (i.e. delayed by only one round trip time (RTT)). This method allows faster convergence compared to current routing protocol.

Pages 19-20, bridging paragraph:

Before explaining the concept of path ~~characterisation~~ characterization metrics further, it should be noted that while the above path ~~characterisation~~ characterization metric appears to correspond in some ways to the TTL value in the IPv4 header, it differs fundamentally from this on account of the fact that the path ~~characterisation~~ characterization information used as feedback is effectively ~~normalised~~ normalized with respect to the receiver rather than the sender. This fundamental difference will be more clearly evident in relation to other embodiments of the invention which may involve path ~~characterisation~~ characterization metrics corresponding to any of a variety of other header values or other characteristics associated with data packets, adapted by applying the above change of reference point to those values or other characteristics in order that they ~~characterise~~ characterize the "downstream path" (i.e. from any node on the path in question) through the network rather than the "upstream path", which is what is ~~characterised~~ characterized by metrics such as the traditional TTL value. A non-

exclusive list of possible candidates which could be used in association with embodiments of the invention follows, together with brief comments on each candidate:

Pages 20-21, bridging paragraph:

Those metrics from the above list that would be necessary and sufficient to operate a simple but complete network service will depend on the type, size and complexity of the network required. The list could include path ~~characterisation~~ characterization metrics corresponding to propagation delay, congestion shadow price and error rate, for example.

Page 21, 2nd full paragraph:

Where the path ~~characterisation~~ characterization metric corresponding to the TTL value or hop-count would in general be decremented in relation to each hop traversed, other mathematical functions may be appropriate in relation to other metrics. Typical ways that the above metrics could be combined between all downstream nodes include the following:

Page 21, penultimate and final paragraphs:

Each path ~~characterisation~~ characterization metric m is represented by a header value h. The header value would in general be combined across all the nodes on a

path with the most useful function, for instance, one of the functions listed above. Logical AND() may be the most appropriate for "Downstream Service Availability", Min() for "Available Capacity", Combinatorial product() for "Congestion shadow price", Difference for "Unloaded delay", etc.

With reference to Figure 4, a graph is shown illustrating the use of a path ~~characterisation~~ characterization metric based on the Explicit Congestion Level (ECL), by way of example.

Page 22, 1st–3rd paragraphs:

A path across a network consisting of a sequence of nodes (v_0, v_1, . . . v_i, . . . v_n) is represented, with source v_0 and destination v_n. Rather than a single bit to notify congestion, a metric "m", used in an Explicit Congestion Level (ECL) multibit field "h" in the network layer header of all data packets is used in this example. This field should be wide enough to represent a reasonable number of discrete values, both positive and negative. Value m_i represents the value of the field before processing by the i^{th} node. It is updated at each node according to a **combining function** *f(.)* common to all the nodes $h_{i+1}(t) = f(h_{i+1}(t), m_i(t))$ where $m_i(t)$ is the local contribution to the end-to-end path characterization metric. $m_i(t)$ may, for instance, be a known value relative to the single downstream link ([[eg]] e.g. the unloaded delay), or reflect some dynamic condition of

the node ([[eg]] e.g. the local congestion level could be given as the dropping probability of the RED algorithm).

A reference value h_z is defined for each metric as the target for the header field at the destination. In Figure 4, $h_z = 0$ for simplicity. Considering now the first of a "flow" of packets (step (1), circled, in Figure 4), the sender or provider (22 in [[Fig.2]] Fig. 2) should estimate an initial value for the ECL, h_0, to place in the packet and store this value. After transmission over the path, the ECL arriving at the destination will be h_n.

The receiver (26 in [[Fig.2]] Fig. 2) then feeds h_n back to the sender using a relevant end-to-end protocol above the network layer (step (2), circled). When this feedback arrives at the sender, any discrepancy between h_n and h_z will require the sender to adjust the initial value it sets the header field to in the data packets it sends. The constraint for h_n to reach h_z at the destination gives the definition of the **source ~~initialisation~~ initialization function** *g(.)* so that $h_0(t+T) = g(h_n(t), h_z, f(.))$. Note that this adjustment occurs one round-trip time after the packet last acknowledged was sent.

Page 23, Table 1:

Combining function		Difference	Combinatorial product	Combinatorial quotient
$h_{i+1}(t) =$		$h_i(t) - m_i(t)$	$1-(1-h_i(t))\cdot(1-m_i(t))$	$1-\frac{(1-h_i(t))}{(1-m_i(t))}$
Source ~~initialisation~~ initialization function	$h_0(t+T) =$	$h_o(t) - h_n(t) + h_z$	$1-\frac{(1-h_0(t))}{(1-h_n(t))}\cdot(1-h_z)$	$1-(1-h_0(t))\cdot(1-h_n(t))$
Downstream path metric extraction function	$\tilde{m}_i(t) =$	$h_i(t) - h_z$	$1-\frac{(1-h_z)}{(1-h_i(t))}$	$h_i(t)$
Metric for which it is useful		Unloaded delay	Congestion	

Page 24, 1st–3rd full paragraphs:

In view of the above, further important advantages of using path ~~characterisation~~ characterization metrics such as the above in the above manner will now be explained with reference to congestion charges, and incentives to act in good faith when providing network status information to other parts of the network. These are as follows:

1) Correct reaction to congestion previously depended on all end-nodes voluntarily complying with standard algorithms. Solutions have been developed in which a price is applied to Explicit Congestion Notification data in packets, giving an incentive to behave responsibly. However, these solutions rely on charging the destination, and expecting it to have a trust relationship with the source in order to

encourage correct source behavior ~~behaviour~~. This has opened the possibility of destinations being subjected to malicious attacks from sources that could force their "victims" to pay congestion charges outside their control. Embodiments of the present invention allow sources to be charged directly for congestion on the downstream path, because information indicative of this is available at the interface between the source and its provider, rather than only at the destination. This also gives the correct incentives and local up-to-date information for inter-connect congestion charging and routing. Currently, each receiving network would have to pay its immediately upstream network in proportion to the number of packets with the congestion experienced code point set in the ECN field. But a downstream network has upstream congestion information but cannot choose who routes to it, and the upstream network doesn't have downstream congestion information but can choose to whom it routes. So downstream networks would have to pay congestion charges to upstream networks whether they would have chosen to have received traffic from them or not.

2) When starting a new flow over a new path, embodiments of the invention provide the correct incentives to proceed cautiously until sufficient feedback has been received. Currently, Internet Protocols require voluntary compliance with congestion control ~~initialisation~~ initialization algorithms in case the path to be used is close to or already in a state of congestion. Such controls lead to conservative behavior ~~behaviour~~, wasting transfer time when a path is in fact nowhere near being congested, which will

become a considerable problem in the future if most objects transfers are complete before feedback from the first packet has arrived. Such controls are also open to abuse, with nodes having an incentive to ignore them for "selfish" reasons. Embodiments of the present invention allow for the risk of lack of knowledge of a path's state to be reflected in the shadow price charged, which may either be ~~realised~~ realized as an actual congestion charge, or as a ~~deprioritisation~~ deprioritization of the traffic carrying the higher shadow prices. It also allows for the correct incentives to be given for intermediate nodes to aggregate numerous flows, each of which separately have no knowledge of the path state, but which can be treated collectively to learn the likely path state for a new flow from the path state recently learned from an old flow.

Page 25, 1st-2nd full paragraphs:

At this stage we can highlight an important point about the congestion level reported in the initial packets in a flow sent without the benefit of any feedback. Although we have already recommended that these values should be flagged as guesses, we still recommend that they should be treated individually just like any other packets. That is, if their downstream congestion level m_{zi} consistently drops below zero, a policing system should ~~penalise~~ penalize (drop) them irrespective of their 'guess' status. So the sender will have to overstate the initial shadow price m_{z0} to ensure such packets have contingency to travel the full length of their unknown path. But the over-

stated shadow price they carry m_{zi} should entitle them to a lesser share of any congested resources, assuming it is higher than other packets of the same class. This effectively enforces a ~~behaviour~~ behavior like the slow start phase of TCP until the path has been correctly characterized ~~characterised~~. Such a harsh regime ensures that the risk of entering an unknown path is borne by the new flow, rather than spread across other flows it encounters.

3) Because knowledge of the downstream path can be available in the network layer header information of downstream data traversing the network, intermediate nodes can use it in order to act as a congestion control proxy for the provider. A specific differentiated services gateway has been invented, which can selectively ~~deprioritise~~ deprioritize and eventually drop the traffic most likely to experience (and therefore cause) congestion on its active downstream paths. Previously, the information required by a proxy was in feedback data passing end to end upstream from destination to source, often on a different path from the downstream flow. Hence proxies found it difficult to access this information, because there was no guarantee they were even on the path of the data. Where such proxies were in use, they were also required to understand all possible higher layer feedback protocols, effectively constraining the introduction of new protocols. Embodiments of the present invention allow for relevant information to be kept at the right layer, in the right direction and therefore on the right path.

Page 26, 1st–2nd paragraphs:

The above discussion relates in general to how embodiments of the present invention allow for solutions of the first of the two general problems set out above, relating to the provision of information ~~characterising~~ characterizing the downstream path to be made available to every node. An explanation will now be given as to how embodiments of the present invention allow for the solution of the second of the two general problems set out above, namely how to proof this information against falsification.

Proofing path ~~characterising~~ characterizing information against falsification

For the following, the explanation will be given with reference to a specific metric, namely a path ~~characterisation~~ characterization metric based on a Congestion Notification field of a future network protocol. In relation to this explanation, it will be assumed that congested nodes decrement the value of the metric by a value indicative of their current level of congestion. A system can be foreseen in which each node sending data along a path (noting that all intermediate nodes, when forwarding data, act in effect both as senders and receivers) pays for the level of congestion it forwards on in sent traffic, and each receiver is paid for the level of congestion in the traffic it receives (except the ultimate receiver – see later). In such a system, intermediate nodes may collect revenue according to the level that they decrement the congestion field in each packet as "congestion charges", and they have an incentive to route packets along the

least congested and hence cheapest downstream path. Each intermediate node may also run a policing algorithm that probabilistically drops packets if their congestion level has decremented below zero, zero being the agreed target level in this case. Dropping algorithms will be discussed later.

Page 27, 4th paragraph:

- Receivers have no incentive to under-declare congestion in their feedback, as this will cause future traffic to be dropped before it reaches them.

Page 29, top center of page: delete "CLAIMS" and insert the following heading:

WHAT IS CLAIMED IS: